

Accelerating the Development of Software-Defined Network Optimization Applications Using SOL

Victor Heorhiadi
UNC Chapel Hill
victor@cs.unc.edu

Michael K. Reiter
UNC Chapel Hill
reiter@cs.unc.edu

Vyas Sekar
Carnegie Mellon University
vsekar@andrew.cmu.edu

ABSTRACT

Software-defined networking (SDN) can enable diverse network management applications such as traffic engineering, service chaining, network function outsourcing, and topology reconfiguration. Realizing the benefits of SDN for these applications, however, entails addressing complex network optimizations that are central to these problems. Unfortunately, such optimization problems require significant manual effort and expertise to express and non-trivial computation and/or carefully crafted heuristics to solve. Our vision is to simplify the deployment of SDN applications using *general* high-level abstractions for capturing optimization requirements from which we can *efficiently* generate optimal solutions. To this end, we present SOL, a framework that demonstrates that it is indeed possible to simultaneously achieve generality and efficiency. The insight underlying SOL is that SDN applications can be recast within a unifying *path-based* optimization abstraction, from which it efficiently generates near-optimal solutions, and device configurations to implement those solutions.

We illustrate the generality of SOL by prototyping diverse and new applications. We show that SOL simplifies the development of SDN-based network optimization applications and provides comparable or better scalability than custom optimization solutions.

1. INTRODUCTION

A key motivation for software-defined networking (SDN) is that it is an enabler for network management applications that would otherwise be difficult to realize using existing control-plane mechanisms. In particular, the consolidation of control logic enables operators to systematically implement network configuration for traffic engineering (e.g., [31]), service chaining (e.g., [30]), energy efficiency (e.g., [16]), network function virtualization (NFV) (e.g., [13]), and cloud offloading (e.g., [35]), among others.

While this body of work has been instrumental in demonstrating the potential benefits of SDN, realizing these benefits requires significant effort. In particular, at the core of many SDN applications are custom optimization problems to tackle various constraints and requirements that arise in practice. (We elaborate on this in §2.) For instance, an SDN application might need to account for limited TCAM, link capacities, or middlebox capacities, among other considerations. Developing such optimization formulations involves a non-trivial learning curve, a careful understanding of both

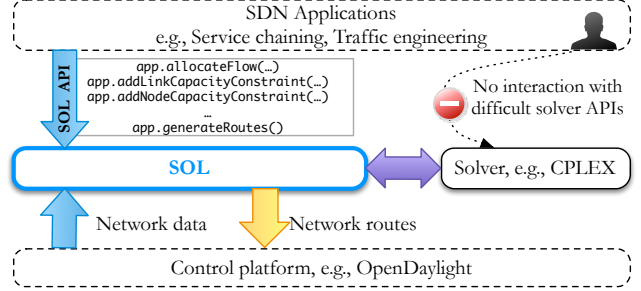


Figure 1: Overview of SOL: Developers use the SOL high-level APIs to specify optimization goals and constraints. SOL generates near-optimal solutions and produces device configurations that are input to the SDN control platform.

theoretical and practical issues, and considerable manual effort. Furthermore, when the resulting optimization problems are intractable to solve with state-of-the-art solvers (e.g., CPLEX or Gurobi), heuristic algorithms must be crafted to ensure that new device configurations can be generated on the timescales demanded by the application as relevant inputs (e.g., traffic matrix entries) change (e.g., [16, 25]). Finally, without a common framework for representing network optimization tasks, it is difficult to reuse key ideas from one application or to combine useful features from multiple applications into a custom new application. As such, many efforts effectively reinvent common building blocks, e.g., ensuring that the rules output by the optimizations can fit inside the TCAM, or generating volume-aware load-balancing rules that also maintain flow affinity.

Our goal in this work is to raise the level of abstraction for writing SDN applications by providing a flexible optimization framework to reduce the development time to build such applications atop SDN. In doing so, we seek to minimize the “pain point” for each network optimization application. To this end, here we introduce SOL, a framework that enables SDN application developers to express new application goals and constraints in a high-level language, and then generates compliant configurations that can be deployed to SDN control platforms directly (see Fig. 1). Conceptually, SOL can be viewed as an intermediate layer that sits between the SDN optimization applications and the actual control platform. Application developers who want to develop new network optimization capabilities write their optimization formulations using the SOL API.

There are two requirements for a framework like SOL:

- **Generality:** First, we want SOL to be *general*; i.e., capa-

ble of expressing the requirements of optimization problems across the full spectrum of SDN applications (e.g., traffic engineering, policy steering, load balancing, and topology management). This entails abstractions that bridge the gap between SDN practitioners and mathematical optimization tools.

- **Efficiency:** Second, SOL should be *efficient*, generating compliant and optimal (or near-optimal) configurations on a timescale that is responsive to application needs as new inputs (e.g., traffic-matrix changes) become known.

Given the diversity of the application requirements and the trajectory of prior work (e.g., [30, 20, 19, 16, 25, 13, 7, 37, 31, 17]), generality and efficiency appear individually difficult, let alone achieving both simultaneously. Our contribution in this paper is (perhaps surprisingly) to show that this is indeed possible, and that generality does not need to come at the expense of efficiency.

Our key insight to achieve *generality* is that SDN applications and their associated optimization problems can be expressed as *path-based* optimization problems. It is very natural for application developers to think in terms of the paths that packets traverse through the network. Moreover, it becomes easy to express key policy requirements in terms of those paths. For example, we can specify service chaining requirements (e.g., each path includes a firewall and intrusion-detection system (IDS), in that order) or redundancy (e.g., each includes two intrusion-prevention systems (IPS), in case one fails open). Finally, it is easy to model device resource consumption, such as routing table and TCAM space, based on the traffic carried on paths that traverse that device.

The natural question then is if this generality sacrifices efficiency. Indeed, if implemented naively, optimization problems expressed over the paths that traffic might travel will introduce efficiency challenges since the number of paths grows exponentially with the network size. Our insight, however, that it is not necessary to consider them all. Specifically, we show that by combining infrequent, offline preprocessing with simple, online *path-selection algorithms* (e.g., shortest paths or random paths), SOL achieves near-optimal solutions in practice for all applications we considered. Moreover, SOL is typically far more efficient than solving the optimization problems originally used to express these applications' requirements.

We have implemented SOL as a Python-based library that interfaces with OpenDaylight (§7). We have also prototyped numerous SDN applications in SOL, including SIMPLE [30], ElasticTree [16], Panopticon [25], and others of our own design (§6). We are planning an open-source release of SOL, and these and other applications coded in SOL, in the near future. Our evaluations on a range of topologies show that: 1) SOL outperforms several of these applications' native optimization algorithms; 2) SOL scales better than other network management tools like Merlin [36]; and 3) SOL substantially reduces the effort re-

quired (e.g., in terms of lines of code) for implementing new SDN applications.

2. BACKGROUND AND MOTIVATION

In this section, we describe representative network applications that could benefit from a framework such as SOL. We highlight the need for careful formulation and algorithm development involved in prior efforts, as well the diversity of requirements they entail.

2.1 Traffic engineering

Traffic engineering (TE) is a canonical network management application that is one of the early driving applications for SDN [20, 19]. Fig. 2 shows an example where classes C1 and C2 need to be routed completely while minimizing the load on the most heavily loaded link. At a high-level, the inputs to a TE application are the traffic demands (e.g., the traffic matrix between WAN sites) and a specification of the traffic *classes* and priorities. The TE application takes into account the network topology and link capacities to determine how to route each class to achieve network-wide objectives; e.g., simple objectives like minimize maximum congestion [10] or weighted max-min fairness [20, 19].




Figure 2: Traffic engineering applications

Challenges: While simple goals like minimizing link congestion can be compactly represented and solved via max-flow formulations [1], the expressivity and efficiency quickly breaks down for more complex objectives such as weighted max-min fairness [20, 19]. In fact, a significant part of the technical contribution in SWAN is the development of an algorithm for max-min fairness [19]. Similarly, B4 [20] has implemented several theoretical optimizations for scalability [8]. When max-flow like formulations fail, designers invariably revert to “low-level” first-principles techniques such as linear programs (LP) or combinatorial algorithms. Neither is desirable; using and debugging with LP solvers is painful as they expose a very low-level interface, and combinatorial algorithms require significant theoretical expertise. Finally, translating the algorithm output into actual routing rules requires care to realize the optimization benefits; e.g., the rules must be volume-aware to realize the benefits of the TE policy [38].

2.2 Service chaining

Networks today rely on a wide variety of third-party appliances or middleboxes for performance, security, and policy compliance capabilities such as intrusion detection and

prevention systems, firewalls, application-layer proxies, and load balancers (e.g., [35]). The goal of service chaining is to ensure that each class of traffic is routed through the desired sequence of network functions to meet the policy requirements. For example, in Fig. 3, class C1 is required to traverse a firewall and proxy, in that order. Such policy routing rules must be suitably encoded within the available TCAM on SDN switches [30]. Since some middleboxes implement expensive computations, they can get easily overloaded and thus operators would like to balance the load on these appliances [30, 14]. Thus, the key inputs to such applications are the different traffic policy classes, their service chaining requirements, and the available middlebox resources in the network. The application then sets up the forwarding rules such that the policy requirements for service chaining are met while respecting the switch TCAM and middlebox capacities. Finally, as a practical constraint, many of these middleboxes are stateful, so the load-balancing rules that are installed must ensure flow affinity.




Figure 3: Service chaining applications

Challenges: Service chaining introduces more complex requirements when compared to TE applications. First, modeling the consumption of switch TCAM introduces “discreteness” into the optimization problem, which impacts scalability [30]. Second, expressing such service processing requirements falls outside the scope of existing network flow abstractions [7]. Third, service chaining also highlights the complexity of combining different network requirements; e.g., reasoning about the interaction between the load balancing algorithm and the switch TCAM constraints is quite non-trivial and introduces circular dependencies [21]. Existing service chaining efforts have to develop custom heuristics [6] or need new theoretical extensions [7]. Finally, ensuring requirements like flow affinity is tricky due to routing asymmetry [18] or due to potential rule conflicts [17] and thus require non-trivial care.

2.3 Flexible topology management

SDN can enable new topology modification capabilities that would otherwise be difficult, if not impossible, to implement with existing control plane techniques. For instance, prior work such as ElasticTree [16] and Response [37] has used SDN to dynamically switch on/off network links and nodes to make datacenters more energy efficient. In Fig. 4, these applications might shut down node N3 during periods of low utilization, if routing classes C1 and C2 via

N4 does not significantly impact end-to-end performance. Topology reconfiguration is especially feasible in rich data-center topologies with multiple paths between every source and destination. Such applications take as input the demand matrix (similar to the TE task) and then compute both which nodes and links should be active and traffic-engineered routes to ensure some performance service-level agreements.




Figure 4: Topology reconfiguration applications

Challenges: The on-off requirement on the switches/links once again introduces discrete constraints, yielding integer-linear optimizations that are theoretically intractable and difficult to directly express using common max-flow like abstractions. As prior work shows, solving such a problem requires significant computation even on small topologies and thus we may have to develop new heuristic solving strategies; e.g., ElasticTree uses a greedy bin-packing algorithm [16].

2.4 Network function virtualization

One of the key advantages of SDN is that it decouples the logical specification of network requirements from their physical realization. Prior work has leveraged SDN capabilities to offload or outsource such expensive functions to leverage clusters or clouds [35, 15, 31]. This is especially useful in the context of complex and expensive deep-packet-inspection services that networks need today [17]. The key decision here is to decide how much of the processing on each path to offload to the remote datacenter — e.g., in Fig. 5, how much of class C1 traffic should be routed to the datacenter between N4 and N5 for IPS processing, versus processing it at N3. Offloading can increase user-perceived latency and impose additional load on network links. Moreover, some active functions (e.g., WAN optimizers or IPS) may induce changes to the observed traffic volumes due to their actions. Thus, optimizing such offloading must take into account the congestion that might be introduced, as well as latency impact and any traffic volume changes induced by such outsourced functions. Generalizing this further, we can also envision novel *elastic scaling* opportunities where the number of middleboxes required can even be scaled on demand [13, 28, 5].

Challenges: These kind of offloading and elastic scaling opportunities introduce new dimensions to optimization that are painful to capture. For instance, with offloading we are effectively rerouting the traffic and thus we need to carefully model the impact on TE objectives and link loads and on




Figure 5: Offloading network functions

downstream nodes. If done naively, this can introduce non-linear dependencies since the actions of downstream nodes depend on control decisions made upstream. To complicate matters further, many of these network functions may actively change the traffic volumes (e.g., compression for redundancy elimination or drops by IPS). Once again, if not done carefully, this introduces non-linear dependencies in the optimization. Finally, elastic scaling once again introduces a discrete aspect to the problem similar to the topology modification application above and further makes the problem intractable.

2.5 Motivation for SOL

Drawing on the above discussion (and our own experience) we summarize a few key considerations:

- Network applications have diverse and complex optimization requirements; e.g., service chaining requires us to reason about “valid” paths and topology modification needs to enable/disable nodes.
- Designers of these applications have to spend significant effort in expressing and debugging these problems using low-level optimization libraries.
- It can take non-trivial expertise to ensure that the problems can be solved fast enough to be relevant for operational timescales, e.g., recomputing TE every few minutes or periodically solving the large ILPs characteristic of topology reconfiguration applications (e.g., [16]).

We argue that if we had better general abstractions for expressing these applications and tools for efficiently solving these problems, the time and effort spent in these aforementioned efforts (and future SDN applications) can be dramatically reduced. Based on the diversity of the requirements we see above and the non-trivial effort¹ that these prior efforts have spent in optimization algorithm development, such a goal may seem elusive. Fortunately, as we show in the rest of the paper, this goal can be realized.

3. SOL OVERVIEW

Our overarching vision in developing SOL is to raise the level of abstraction in developing new SDN applications and specifically to eliminate some of the black art in develop-

¹While it is hard to directly estimate the number of human-hours spent, using the real estate in these papers devoted to the optimization component as a rough proxy suggests non-trivial effort.

ing SDN-based optimizations, making them more accessible for deployment by network managers. SOL takes in as inputs the network topology and traffic patterns and the optimization requirements specified by clients in the SOL API. It will then translate these into low-level constraints in the language of existing optimization solvers such as CPLEX or Gurobi. Finally, SOL interfaces with existing SDN control platforms such as OpenDaylight to install the desired forwarding rules on the SDN switches and also install other middlebox-specific configuration parameters. SOL does not require modifications to the existing control or data plane components of the network. Our vision for SOL stands in stark contrast to the state of affairs today, in which a developer faces programming a new SDN optimization either directly for a generic and low-level optimization solver such as CPLEX or using a typically heuristic algorithm designed by hand, after which she must translate the decision variables of the optimization to device configurations.

Path abstraction: For SOL to be useful and robust, we need a unifying abstraction that can capture the requirements of diverse classes of SDN network optimization applications described in the previous section. SOL is built using *paths* through a network as a core abstraction for expressing network optimization problems. This is contrary to how many optimizations are formulated in the literature — using a more standard edge-centric approach [1]. In our experience, however, an edge-centric approach forces complexity when presented with additional requirements, especially ones that attempt to capture path properties [25, 16].

In contrast, path-based formulations capture these requirements more naturally. For instance, much of the complexity in the modeling for policy steering or offloading applications from §2 is in capturing the path properties that need to be satisfied. With a path-based abstraction, we can simply define predicates that specify valid paths — e.g., those that include certain waypoints or that avoid a certain node (to anticipate that node’s failure). In addition, we can model path-based resource use with ease. For example, usage of TCAM space in a switch corresponds to a traffic-carrying path traversing that switch (and thus a rule to accommodate that path). Without the path abstraction, modeling such constraints is difficult (cf., [30]). Finally, expressing constraints on nodes and edges requires little change and does not introduce increased difficulty.

Tractability: In a pure flow-routing scenario, an edge-based formulation admits simple algorithms that guarantee polynomial-time execution. Path-based formulations, on the other hand, are often dismissed because of their inefficient appearance — after all, in the worst case, number of paths in the network is exponential with respect to network size — or due to the complexity of algorithms to solve path based formulations (column-generation, decompositions, etc. [1]). However, in many practical scenarios, the number of valid paths (as defined by the application) is likely to be significantly smaller. Furthermore, multipath routing can pro-

vide only so much network diversity before its value diminishes [26]. So, the set of paths considered need not be large.

SOL leverages an off-line path generation step to determine valid paths (step 1 of Fig. 6). Since for most applications, the set of valid paths is fairly static and does not need to be recomputed every time the optimization is run, this step is infrequent. Next, SOL *selects* a subset of these paths (step 2) using a selection strategy (see §5) and runs the optimization using only the selected paths as those available (step 3), to ensure that the optimization completes quickly. We show in §8 that this strategy still permits inclusion of sufficiently many paths for the optimization to converge on a (near) optimal value. So, while the efficiency of path-based optimization is a valid theoretical concern, in practice we show that there are practical heuristics to address this issue.

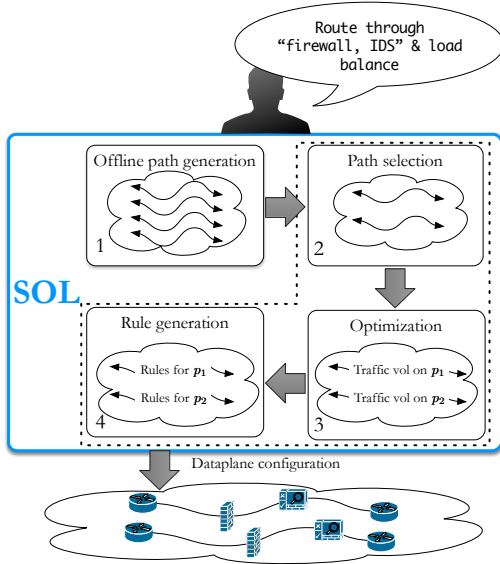


Figure 6: SOL architecture, overview of the workflow

Generating device configurations: SOL directly translates the decision variables of a SOL optimization to network device configurations to, for example, implement appropriate flow routing (step 4 of Fig. 6). The algorithm utilized in SOL to perform this translation is based on that in previous work [38, 17], but is critically enabled by the path-based formulations in SOL — one more advantage of the path abstraction that we employ. In contrast, the original optimization formulations of some applications we consider require the development of an additional custom algorithm to map their decision variables to device configurations.

4. SOL DETAILED DESIGN

In this section, we present the detailed design of SOL. We focus on the high-level API that the SDN application developer would use to express applications via SOL. For completeness, however, we present the formal basis that SOL uses *internally* for each such template.

Note, however, that the developer “thinks” in terms of the

`nodes`: Set of all nodes, part of the topology

`links`: Set of all links, part of the topology

`classes`: Set of all traffic classes

`paths(c)`: Paths available for class $c \in \text{classes}$; output by path-selection stage (§5)

Figure 7: Network data input

	Var.	API	Description
Decision	$x_{c,p}$	<code>xp(c, p)</code>	Fraction of class- c flows allocated to path $p \in \text{paths}(c)$; non-integer
	b_p	<code>bp(p)</code>	Is path p used; binary
	b_v	<code>bn(v)</code>	Is node v used; binary
	b_l	<code>be(l)</code>	Is link l used; binary
	$capvar_v^r$	<code>nc(v, r)</code>	Capacity allocated for resource r at node v ; non-integer
Derived	a_c	<code>al(c)</code>	Fraction of c ’s “demand” routed; non-integer
	$load_l^r$	<code>el(l, r)</code>	Amount of resource r consumed by flows routed over link l ; non-integer
	$load_v^r$	<code>nl(v, r)</code>	Amount of resource r consumed by flows routed via node v ; non-integer

Figure 8: Variables internal to the optimization

high-level API rather than low-level details of dealing with the solver-level variables, how paths are identified, etc. A developer begins a SOL-based optimization by instantiating a new `opt` object via the `getOptimization` function and then proceeds to build the optimization using the *constraint templates*, which we explain below.²

4.1 Preliminaries

Data inputs: There are two basic data inputs that the developer needs to provide to any network optimization. First, the network topology is a required input, specified as a graph with `nodes` and `links`. It also contains metadata of node/edge types or properties; e.g., nodes can have designated functions like “switch” or “middlebox”. Second, SOL needs a specification of *traffic classes*, where each class c has associated ingress and egress nodes and some expected traffic volume. Each class can (optionally) be associated with a specification of the “processing” required for traffic in this class, e.g., service chaining. Finally, to each traffic class c is associated a set `paths(c)` available to route flows in class c ; `paths(c)` is output by a path-selection preprocessing step described in §5.

Internal variables: SOL internally defines a set of variables summarized in Fig. 8. We reiterate that the developer does not need to reason about these variables and uses a high-level mental model as discussed earlier.³ There are two main kinds of variables:

- *Decision variables* that identify key optimization con-

²Due to space limits we focus on salient aspects of the API and refer readers to our (anonymized) manual [3].

³We also expose low-level APIs (§4.6) for advanced users.

trol decisions. The most fundamental decision variable is $x_{c,p}$, which captures traffic routing decisions and denotes the fraction of flow for a traffic class c that path $p \in \text{paths}(c)$ carries. This variable is central to various types of resource management applications as we will see later. To capture topological requirements (e.g., §2.3), we introduce three binary decision variables b_p , b_v , and b_l that denote whether each path, node or link (respectively) is enabled ($= 1$) or disabled ($= 0$). The variable capvar_v^r is the SOL-given allocation of resource- r to node v .

- *Derived variables* are functions defined over the above decision variables that serve as convenient “shorthands”. a_c denotes the total fraction of flow for class c that is carried by all paths. The load variables load_v^r and load_l^r model the consumption of resource r on node v and link l , respectively.

The (low-level) API calls in Fig. 8 return the name of the corresponding variable, which can be used to access its value in a public map of variable names to values.

Given these preliminaries, next we describe the key building blocks of a SOL-based network optimization program.

4.2 Routing requirements

Routing constraints control the allocation of flow in the network. `addAllocateFlowConstraint` creates the necessary structure for routing the traffic through a set of paths for each traffic class. Some network applications try to satisfy as much of their flow demands as possible (e.g., max-flow) while others (e.g., TE) want to “saturate” demands. For example, a developer of a TE application (§2.1) would like to route all traffic through the network, and thus she would add the following high-level routing constraint templates to her empty `opt`:

```
opt.addAllocateFlowConstraint()
opt.addRouteAllConstraint()
```

A developer writing a simple max-flow, however, would only need `addAllocateFlowConstraint` since there is no requirement on saturating demands in that case.

The `addEnforceSinglePath(C)` constraint forces a single flow-carrying path per class $c \in C$, preventing flow-splitting and multipath routing.

Internals: Formally, `addAllocateFlowConstraint` ensures that the total traffic flow across all chosen paths for the class c matches the variable a_c .

$$\forall c \in \text{classes} : \sum_{p \in \text{paths}(c)} x_{c,p} = a_c$$

Similarly, `addRouteAllConstraint` implies:

$$\forall c \in \text{classes} : a_c = 1$$

Due to space limitations, we do not provide the formal basis for `addEnforceSinglePath`.

4.3 Resource capacity constraints

`linkCapFn(l, c, p, r)`: Amount of resource type r consumed if all class- c traffic is allocated to path $p \ni l$ for link l

`nodeCapFn(v, c, p, r)`: Amount of resource r consumed if all class- c traffic is allocated to path $p \ni v$ for node v

`nodeBudgetFn(v)`: Cost of using node v ; required with `addBudgetConstraint`

`routingCostFn(p)`: Cost of routing along path p ; required with `minRoutingCost`

`predicate(p)`: Determine whether any given path is valid by returning `True` or `False`

Figure 10: Customizable functions

As we saw in §2, SDN optimizations have to deal with a variety of capacity constraints dealing with network resources, such as link bandwidth, switch rule capacities, middlebox CPU and memory capacities. SOL provides users the flexibility to write custom resource management logic within a given network structure. For this purpose we create several types of functions, depicted in Fig. 10. These functions allow the user to compute the “cost” of routing traffic through a link, a node, or a given path. SOL provides multiple implementations of these for common tasks, but allows the user to specify their own logic, as well, as we will show later (§6).

For example, `addLinkCapacityConstraint` deals with constraints on the network links, and `addNodeCapacityConstraint` incorporates constraints on the nodes. For our example, let us add a constraint that limits link usage. For this, we need a resource that we are constraining (in this case, *bandwidth*), a mapping of links to their capacities,⁴ and optionally, a way to compute the cost of traffic on a link.

```
opt.addLinkCapacityConstraint
    ('bandwidth',
     {(1,2): 10**7, (2,3): 10**7},
     defaultLinkFunction)
```

This indicates that bandwidth should not exceed 10 Mbps for links 1-2 and 2-3. Note that the default function is purely for illustration; the developer can write her own `linkCapFn` (recall Fig. 10).

`addNodeCapacityPerPathConstraint` generates constraints on the nodes that do not depend on the traffic, but rather on the routing path. That is, the cost of routing at a node does not depend on the volume or type of traffic being routed; it depends on the path and its properties. The best example of such usage is accounting for the limited rule space on a network switch (e.g., §2.2). If a path is “active”, the rule must be installed on each switch to support the path.

Internals: `addLinkCapacityConstraint` and `addNodeCapacityConstraint` rely on `linkCapFn`

⁴When capacities are unknown a priori and should be allocated by the optimization itself, a capacity of TBA (meaning To Be Allocated) can be specified, instead.

Group	Function	Description
Routing ($C \subseteq \text{classes}$)	addAllocateFlowConstraint	Allocate flow in the network
	addRouteAllConstraint	Route all traffic demands
	addEnforceSinglePath(C)	For each $c \in C$, at most one $p \in \text{paths}(c)$ is enabled.
Capacities	addLinkCapacityConstraint($r, \text{lnCap}, \text{linkCapFn}$)	If l is in lnCap , then limit utilization of link resource r on link l to $\text{lnCap}[l]$.
	addNodeCapacityConstraint($r, \text{ndCap}, \text{nodeCapFn}$)	If v is in ndCap , then limit utilization of node resource r on node v to $\text{ndCap}[v]$.
	addNodeCapacityPerPathConstraint($r, \text{ndCap}, \text{nodeCapFn}$)	If v is in ndCap , then limit utilization of node resource r on node v by enabled paths to $\text{ndCap}[v]$.
	addCapacityBudgetConstraint(r, N, totCap)	Limit total type- r resources allocated to nodes in $N \subseteq \text{nodes}$ to totCap . Used when SOL is allocating capacities.
	addRequireAllNodesConstraint(C)	For each $c \in C$ and each $p \in \text{paths}(c)$, p can be enabled iff all nodes on p are enabled.
Topology control ($C \subseteq \text{classes}$)	addRequireSomeNodesConstraint(C)	For each $c \in C$ and each $p \in \text{paths}(c)$, p can be enabled iff some node on p is enabled.
	addRequireAllEdgesConstraint(C)	For each $c \in C$ and each $p \in \text{paths}(c)$, p can be enabled iff all links on p are enabled.
	addPathDisableConstraint(C)	For each $c \in C$ and each $p \in \text{paths}(c)$, p can carry traffic only if it is enabled.
	addBudgetConstraint($\text{nodeBudgetFn}, k$)	Total cost of enabled nodes, as computed using nodeBudgetFn , is at most k .
Objective	setPredefinedObjective(name)	Set one of the predefined functions as the objective (see Fig. 11).

Figure 9: Selected constraint template functions for building optimizations; see Fig. 10 for linkCapFn , nodeCapFn , and nodeBudgetFn

and nodeCapFn , respectively, to compute the cost of using a particular resource at a link or node if all of the class- c traffic was routed to it. Internally, the load is multiplied by the $x_{c,p}$ variable to capture the load accurately, then the load is capped by a user-provided lnCap (ndCap), which is a mapping of links (nodes) to capacities for a given r . (Similar node capacity equations not shown for brevity.)

$\forall l \text{ in } \text{lnCap} :$

$$\text{load}_l^r = \sum_c \sum_{p \in \text{paths}(c): l \in p} x_{c,p} \times \text{linkCapFn}(l, c, p, r)$$

$$\text{load}_l^r \leq \text{lnCap}[l]$$

The `addNodeCapacityPerPathConstraint` functions a bit differently, as it depends on enabled paths:

$\forall v \text{ in } \text{ndCap} :$

$$\text{load}_v^r = \sum_c \sum_{p \in \text{paths}(c): v \in p} b_p \times \text{nodeCapFn}(v, c, p, r)$$

$$\text{load}_v^r \leq \text{capvar}_v^r$$

$$\text{if } \text{ndCap}[v] \neq \text{TBA} \text{ then } \text{capvar}_v^r = \text{ndCap}[v]$$

4.4 Node/link activation constraints

These constraints, when used, allow developers to logically model the act of *enabling* or *disabling* nodes, links, and paths, e.g., for managing energy or other costs (e.g., §2.3). We identify two possible modes of interactions between these topology modifiers and the optimization. The optimization developer can choose the one that is most suitable for their context: `addRequireAllNodesConstraint` captures the property that disabling a node disables all paths that traverse it; and `addRequireSomeNodesConstraint` captures

the property that enabling a node permits any path traversing it to be enabled, as well. The latter version is suitable when, e.g., a node can still route traffic even if its other (middlebox) functionality is disabled, and so a path containing that node is potentially useful as providing middlebox functions if at least one of its nodes is enabled. Naturally there are analogous constraint templates for links, but we omit them here for brevity. A third constraint template, `addPathDisableConstraint`, restricts a path to carry traffic only if it is enabled.

For example, a developer trying to implement the application from §2.3 can model the requirements for shutting off datacenter nodes by adding the `addRequireAllNodesConstraint` and `addPathDisableConstraint` templates:

```
opt.addRequireAllNodesConstraint
    (trafficClasses)
opt.addPathDisableConstraint
    (trafficClasses)
```

Other efficiency considerations may enforce a budget on the number of enabled nodes, to model constraints on total power consumption of switches/middleboxes, cost and budget of installing/upgrading particular switches, etc. These are captured via the `addBudgetConstraint` constraint.

Internals: Internally, these topology modification templates are achieved using the binary variables we introduced earlier. Specifically, the above requirements can be formalized as follows:

$$\begin{array}{ll} \text{addRequireAllNodesConstraint} & \forall p \in \text{paths}(c) : \\ \text{addRequireSomeNodesConstraint} & \forall v \in p : b_p \leq b_v \\ \text{addPathDisableConstraint} & b_p \leq \sum_{v \in p} b_v \\ & x_{c,p} \leq b_p \end{array}$$

maxAllFlow	maximize $\sum_{c \in \text{classes}} a_c$
minMaxNodeLoad (r)	minimize $\max_{v \in \text{nodes}} \text{load}_v^r$
minMaxLinkLoad (r)	minimize $\max_{l \in \text{links}} \text{load}_l^r$
minRoutingCost	$\sum_{c,p} \text{routingCostFn}(p) \times x_{c,p}$

Figure 11: Common objective functions

Naturally, similar constraints are constructed for links. Note that `addPathDisableConstraint` is crucial to the correctness of the optimization in that it enforces that no traffic traverses a disabled path. For brevity, we do not provide the formal basis for `addBudgetConstraint`.

4.5 Specifying network objectives

The goal of SDN applications is eventually to optimize some network-wide objective, e.g., maximizing the network throughput, load-balancing, or minimizing total traffic footprint. Fig. 11 lists the most common objective functions, drawing on the applications considered in §2. For instance, the developer of a TE application may want to implement the objective of minimizing the maximum link load and thus add the following code snippet:

```
opt.setPredefinedObjective
(minMaxLinkLoad)
```

Other developers (e.g., §2.4) may want to minimize the total routing cost and invoke a `minRoutingCost` objective. This objective is parameterized via `routingCostFn(p)`; i.e., developers can plugin their own cost metrics such as number of hops or link weights. As shown, we also provide a range of natural load-balancing templates.

4.6 Advanced users and low-level interface

As we will see in §6, the SOL API is general and expressive enough to capture the diverse requirements of the broad spectrum of applications discussed in §2. That said, as part of the SOL design we also expose a low-level API that gives more control to the user by giving access to the SOL internal variables. Advanced users can use this API for further customization, as we will see in §6.

For instance, each API call in Fig. 8 enables the name of the corresponding internal variables to be retrieved. Similarly, using the `defineVar(name, coeffs, lb, ub)` function, the user can create a new variable with name `name`, specify numeric lower and upper bounds (`lb` and `ub`), and equate it to a linear combination of any other existing variables as specified by `coeffs`, a map from variable names to numeric coefficients. This is a useful primitive when specifying complex objectives. Finally, SOL also allows to set a custom objective function that is a linear combination of any existing variables. This is done using the `setObjective(coeffs, dir)` function call, which accepts a mapping `coeffs` of variable names to their coefficients. The binary input `dir` indicates whether the objective should be minimized or maximized.

5. PATH GENERATION AND SELECTION

Given these constraint templates, the remaining question is how we populate the path set `paths(c)` for each traffic class `c` to meet two requirements. First, each $p \in \text{paths}(c)$ should satisfy the desired policy specification for the class `c`. Second, `paths(c)` should contain paths for each class `c` that make the formulation tractable and yet yield near-optimal results. We describe how we address each concern next.

Generation: First, to populate the paths, SOL does an offline enumeration of all simple (i.e., no loops) paths per class.⁵ Given this set, we filter out the paths that do not satisfy the user-defined predicate `predicate` such that `predicate(p) = True` only if p is a “valid” path. (We can generalize this to allow different predicates per class; not shown for brevity.)

In practice, we implement the predicate as a flexible Python callable function rather than constrain ourselves to specific notions of path validity (e.g., regular expressions as in prior work [36]). Using this predicate gives the user flexibility to capture a range of possible requirements. Examples include waypoint enforcement (forcing traffic through a series of middleboxes in order); enforcing redundant processing (e.g., through multiple NIDS, in case one fails open); and limiting network latency by mandating shorter paths.

Selection: Using all valid paths per class may be inefficient as the number of paths grows exponentially with the size of the network; i.e., the LP/ILP SOL generates will have an exponential number of rows and columns making it difficult to solve in reasonable time. SOL provides path selection algorithms that choose a subset of valid paths that are still likely to yield near-optimal results in practice. Specifically, two natural methods work well across the spectrum of applications we have considered: (1) `selectNumber` shortest paths for latency-sensitive applications (`selectStrategy = shortest`) or (2) `selectNumber` random paths for applications involving load balancing (`selectStrategy = random`). SOL is flexible to incorporate other selection strategies; e.g., picking paths with minimal node overlap for fault tolerance.

Developer API: The developer can specify the path predicate and selection strategy when starting the optimization. Note that the developer does not need to be involved in the low-level details of generation and pruning; SOL runs these steps automatically. We also provide APIs for developers to add their own logic for generation and selection; we do not discuss these due to space limitations.

6. EXAMPLES

Next, we show end-to-end examples to highlight the ease of using the SOL APIs to write existing and novel SDN network optimizations. These examples are actual Python code that can be run, not just pseudocode (we refer the reader

⁵This is to simplify the forwarding rules without resorting to tunneling or packet tagging [30].

```

1  SIMPLE_predicate = functools.partial(
    waypointMboxPredicate, order=('fw', 'ids'))
2  def SIMPLE_NodeCapFunc(node, tc, path, resource,
    nodeCaps):
3      if resource == 'cpu' and node in nodeCaps['cpu']:
4          return tc.volFlows * tc.cpuCost / nodeCaps[
            resource][node]
5  capFunc = functools.partial(SIMPLE_NodeCapFunc,
    nodeCaps=nodeCaps)

6
7  def SIMPLE_TCAMFunc(node, tc, path, resource):
8      return 1
9  # Path generation, typically run once in a
    precomputation phase
10 opt = getOptimization()
11 pptc = generatePathsPerTrafficClass(topo,
    trafficClasses, SIMPLE_predicate, 10, 1000,
    functools.partial(useMboxModifier, chainLength
    =2))
12 # Allocate traffic to paths
13 pptc = chooserand(pptc, 5)
14 opt.addDecisionVariables(pptc)
15 opt.addBinaryVariables(pptc, topo, ['path', 'node'])
16 opt.addAllocateFlowConstraint(pptc)
17 opt.addRouteAllConstraint(pptc)
18 opt.addLinkCapacityConstraint(pptc, 'bandwidth',
    linkCaps, defaultLinkFuncNoNormalize)
19 opt.addNodeCapacityConstraint(pptc, 'cpu',
    {node: 1 for node in topo.nodes() if 'fw' or '
    ids' in topo.getServiceTypes(node)}, capFunc)
20 opt.addNodeCapacityPerPathConstraint(pptc, 'tcam',
    nodeCaps['tcam'], SIMPLE_TCAMFunc)
21 opt.setPredefinedObjective('minmaxnodeload', 'cpu')
22 opt.solve()
23 obj = opt.getSolvedObjective()
24 pathFractions = opt.getPathFractions(pptc)
25 c = controller()
26 c.pushRoutes(c.getRoutes(pathFractions))

```

Figure 12: Python code to express SIMPLE [30] in SOL

to the anonymized manual for full code examples [3]). By comparison, the code is significantly higher-level and more readable than the equivalent CPLEX code would be, as it does not need to deal with large numbers of underlying variables and constraints.

Service chaining (§2.2): As a concrete instance of the service chaining example, we consider SIMPLE [30]. SIMPLE involves the following requirements: route all traffic through the network, enforce the service chain (e.g., “firewall followed by IDS”) policy for all traffic, load balance across middleboxes, and do so while respecting CPU, TCAM, and bandwidth requirements. Fig. 12 shows how SIMPLE optimization can be written in ≈ 25 lines of code. This listing assumes that topology and traffic classes have been set up, in the `topo` and `trafficClasses` objects, respectively.

The first part of the figure shows function definitions and the path generation step, which would typically be performed once as a precomputation step. We start by defining a path predicate (line 1) for basic enforcement through middleboxes by using the SOL-provided function with the middlebox order. The next few lines (lines 2-4) show a custom node capacity function to normalize the CPU load between 0 and 1. This computes the processing cost per traffic class (number of flows times CPU cost) normalized by the current node’s capacity. Similarly, the TCAM capacity function captures that each path consumes a single rule per switch (line 7). The user gets the optimization object (line 10),

and generates the paths (line 11), obtaining the “paths per traffic class” (`pptc`) object. The path generation algorithm is parameterized with the custom `SIMPLE_predicate`, a limit on path length of 10 nodes, and a limit on the number of paths per class of 1000. It is also instructed to evaluate every possible use of two middleboxes on a routing path for inclusion as a distinct path in the output.

The remaining lines show what would be executed whenever a new allocation of traffic to paths is desired. Line 13 selects 5 random paths per traffic class; lines 14–20 add the routing and capacity constraints. We use the default link capacity function for bandwidth constraints, and our own functions for CPU and TCAM capacity. Because the CPU capacity function normalizes the load, the capacity of each node is now 1 (line 19). The user then selects a predefined objective to minimize the CPU load (line 21) and calls the solver (line 22). Finally, the program gets the results and interacts with the SDN controller interface to automatically install the rules (line 26).

ElasticTree [16]: Due to space limitations we only show the most important difference between ElasticTree and SIMPLE. First, there is no requirement on paths, and so `nullPredicate` is used for path generation. Also, we use link binary variables (see line 1 below), as well as the node/link activation constraints (lines 2–4). Finally, we must use the low-level API to define power consumption for switches and links (lines 5, 6) and use these variables to define a custom objective function (line 7).

```

1  opt.addBinaryVariables(pptc, topo, ['path', 'node',
    'edge'])
2  opt.addRequireAllNodesConstraint(pptc)
3  opt.addRequireAllEdgesConstraint(pptc)
4  opt.addPathDisableConstraint(pptc)
5  opt.defineVar('SwitchPower', {opt.bn(node):
    switchPower[node] for node in topo.nodes()})
6  opt.defineVar('LinkPower', {opt.be(u, v): linkPower
    [(u, v)] for u, v in topo.links()})
7  opt.setObjectiveCoeff({'SwitchPower': .75, '
    LinkPower': .25}, 'min')

```

New elastic scaling capabilities: Finally, we show SOL can be used for novel SDN applications. Specifically, we consider an elastic NFV setting [13] that places middleboxes in the network and allocates capacities on demand in response to observed demand. There could be additional constraints, such as the total number of such VM locations. As a simple objective, we consider an upper bound on the number of nodes used while still load balancing across virtual middlebox instances. We can easily add other objectives such as minimizing number of VMs. For brevity we highlight only the key parts of building such a novel application.

```

1  predicate = hasMboxPredicate
2  opt.addBinaryVariables(pptc, topo, ['path', 'node'])
3  opt.addNodeCapacityConstraint(pptc, 'cpu', {node: '
    TBA' for node in topo.nodes()}, lambda node, tc,
    path, resource: tc.volFlows * tc.cpuCost)
4  opt.addRequireSomeNodesConstraint(pptc)
5  opt.addPathDisableConstraint(pptc)
6  opt.addBudgetConstraint(topo, lambda node: 1, topo.
    getNumNodes() / 2)

```

```

7  opt.setPredefinedObjective('minmaxnodeload',
    resource='cpu')

```

First, we define a valid path to be one that goes through a middlebox; SOL provides a predicate for that (line 1). The main difference here is the definition of capacities with the TBA value on line 3; this indicates that our optimization must allocate the capacities to the nodes. (SOL ensures that disabled nodes have 0 capacity allocated.) Thus we require at least one enabled node per path (lines 4, 5), limit the number of enabled nodes (line 6), and set the objective (line 7).

7. IMPLEMENTATION

Next, we briefly describe how we prototype the various aspects of SOL.

7.1 Core functionality

Developer interface: We currently provide a Python API for SDN optimization that is an extended version of the interface described in §4.

Solvers: We use CPLEX as our underlying solver and use the Python API that CPLEX already supports. The choice largely reflects our familiarity with the tool, and we can substitute CPLEX with other solvers like Gurobi, which has similar capabilities. SOL also has APIs to exploit solver capabilities to use a previously computed solution and incrementally find a new solution. This approach is typically faster than starting from scratch and so is useful for faster reconfigurations.

Path enumeration: Path generation is an inherently parallelizable process; we simply launch separate Python processes for different traffic classes. We currently support two path selection algorithms: `random` and `shortest`. It is easy to add more algorithms as new applications emerge.

Rule generation and control interface: We use OpenDaylight as the SDN control platform and use its REST API to install the relevant rules. We generate the rules based on the optimization output, using network prefix splitting to implement the fractional responsibilities represented by the $x_{c,p}$ variables. This step is similar to prior works that have mapped fractional processing and forwarding responsibilities to flows in the network (e.g., [38, 17]), and so we do not repeat it here.

7.2 Extensions

We have implemented additional features in SOL that are useful for responding to traffic changes or other events.

Minimizing reconfiguration changes: Networks are in flux during reconfigurations with potential performance or consistency implications, and thus we want to minimize the amount of path churn. We extend the basic APIs to support a new constraint that bounds (or minimizes) the logical distance between the previous solution and the new solution to help minimize the number of flows that have to be assigned

a new route. To this end, SOL also supports path selection that gives priority to previously selected paths.

Reacting to faults: We provide a basic mechanism to react to node or link failures as follows. We build a dependency graph between the set of paths we have chosen during generation (§5) and the specific links, switches, and middlebox nodes that these paths traverse. Given this, we implement the following two-step heuristic. As a first step, we simply rerun the optimization with the old set of selected paths but excluding the specific paths impacted by the failures. If the objective value is much worse, then we fall back to running the selection step again (i.e., selecting `selectNumber` paths from the already generated paths for each traffic class, either `random` or `shortest`) and rerun the optimization. This two-step heuristic works well in the common case because (i) the time to run the optimization is lower than the time for selection, and (ii) we can often find near-optimal solutions even with fewer paths per class (Fig. 15).

8. EVALUATION

In this section we show that SOL is practical and yields optimal solutions for an array of representative applications and topologies, at timescales that are often dramatically better than custom solutions. We specifically evaluate the effect of using SOL to implement three existing SDN applications: ElasticTree [16], SIMPLE [30], and Panopticon [25]. For each application, we implemented the original optimization formulation presented in prior work. Specifically, this is an ILP in the case of ElasticTree, SIMPLE, and Panopticon, and its solution (where we could compute it) is denoted as an “original solution” below. We also implemented in SOL the “new elastic scaling” application from §6, to which we refer as “ElasticScaling” below. We do not have a competing non-SOL implementation for ElasticScaling, however. Finally, in order to draw comparisons to a conceptually similar, recent effort (Merlin [36]) in §8.1, we implemented the example application used in that paper, in SOL.

Our evaluation has three parts. In §8.1, we demonstrate that SOL generates optimal (or very nearly optimal) solutions, but does so orders-of-magnitude faster than their original formulations. In §8.2, we describe the improvements to developer effort (e.g., in lines of code needed) that SOL provided in our implementations. In §8.3, we evaluate the sensitivity of SOL solutions to its path selection parameters.

To perform these evaluations, we chose topologies of various sizes from the TopologyZoo dataset [24]. For ElasticTree, we also constructed FatTree topologies of various sizes [2]; we refer to these as “kX” where X denotes the arity of the FatTree, as defined in prior work. Lacking traffic matrices for these networks, we generated them synthetically: we used a uniform traffic matrix for the “kX” networks and a gravity-based model [34] for the TopologyZoo dataset. Randomly sampled values from the log-normal distribution served as “populations” for the gravity model. In our experience, the choice of traffic matrix does not qualita-

	Abilene	Quest	Geant2012	Bellcanada	Dfn	k4
ElasticTree	N/A	N/A	N/A	N/A	N/A	1.1
Panopticon	1	1	1	1	1	N/A
SIMPLE	1	1	N/A	N/A	N/A	N/A

Table 1: Objective ratio of SOL vs. original solutions




Figure 13: Optimization runtime of SOL and the original formulations; gray regions indicate where original formulation could not be solved within 30 mins

tively impact the results we present. Unless otherwise specified, we use 5 paths per traffic class when running SOL. All times below refer to computation on computers with 2.4GHz cores and 128GB of RAM. In the interest of space, we do not present all measures for all applications, but our reported observations are representative across all applications and topologies considered.

8.1 Optimality and scalability

First, we examine how well SOL’s results match original solutions, which are themselves optimal (by definition). Table 1 shows the “objective ratio” of the SOL solution’s objective value to the original solution’s objective value, for topologies and applications for which we were able to obtain original solutions. A ratio of 1 indicates that SOL obtains an optimal solution, which it does in most of the cases. “N/A” indicates that we could not obtain an original solution within 30 minutes of computation.

SOL solution times are at least one order of magnitude faster than is solving the original formulations, and are often two or even three orders of magnitude faster. Fig. 13 shows run times to find original solutions. Again, the runtime was capped at 30 min (1800 s), after which the execution was aborted. Several original formulations did not complete in that time, such as SIMPLE for topologies Bellcandada and larger, and Panopticon for Ion and larger. The topologies for which original solutions could not be found are indicated in the gray regions in Fig. 13.

Comparison to Merlin: Merlin [36] is a recent work that tackles problems of network resource management similar to SOL. While the goals and formulations of Merlin and SOL are quite different, a comparison highlights the generality of SOL and the power of its path abstraction. Specifically, Merlin uses a more heavyweight optimization that is always an ILP [36] and operates on a graph that is substantially larger




Figure 14: Comparison of optimization runtimes of SOL vs Merlin; gray region indicates where Merlin did not find solution within 30 mins

than the physical network. We implemented the example application taken from the Merlin paper using both SOL and Merlin. Fig. 14 shows that SOL outperforms Merlin by two or more orders of magnitude.

Path selection and generation costs: Path selection times are small, ranging from 0.1 to 3 seconds across topologies. Path selection is preceded by a path generation phase that enumerates the simple paths per class. Path generation is moderately costly for large topologies, e.g., taking <300 s for the largest presented topology, when parallelized to 60 threads. However, we highlight that path generation can be relegated to an offline pre-computation phase that is only performed once.

8.2 Developer benefits

We believe that SOL is a much simpler framework for encoding SDN optimization tasks, versus developing custom solutions by hand for solution by an off-the-shelf solver. In an effort to demonstrate this simplicity somewhat quantitatively, Table 2 shows the number of lines of SOL code in our implementations for the various applications (“SOL lines of code”), and the ratio of the lines of code in our hand-developed scripts⁶ to produce the original formulations to the lines of code for our SOL implementations (“Estimated improvement”). We acknowledge that lines-of-code comparisons are inexact, at best, but we do not know of other ways of comparing “development effort” without conducting user studies. (We are considering this option for future work.)

Name	SOL lines of code	Estimated improvement
ElasticTree	16	21.8×
Panopticon	13	25.7×
SIMPLE	21	18.6×
ElasticScaling	15	N/A

Table 2: Development effort benefits provided by SOL

There are reasons to believe, however, that the improvements indicated in Table 2 are even conservative. First, in our opinion, our scripts for producing original formulations

⁶Hand-coding original formulations for direct consumption by CPLEX would imply repeating that process for each topology, and so we scripted the generation of original formulations in Python as efficiently as we could.

are considerably more complex and delicate than our SOL code. We primarily attribute this difference to needing to account for CPLEX particulars at all; with SOL, these particulars are completely hidden from the developer. Second, SOL translates its optimization results to device configurations, whereas this functionality is not even included in our scripts for producing original formulations. Producing device configurations from original solutions would require designing an algorithm to map the control variables in each formulation to relevant device configurations.

8.3 Sensitivity

SOL solutions require the specification of both a number (`selectNumber`) and type (`shortest` or `random`) of paths to select per traffic class. In this section we quantify how sensitive SOL is to these path selection parameters.

Number of Paths: Fig. 15 shows the runtime and the objective ratio as a function of the number of paths per class for two applications, SIMPLE and Panopticon. Unsurprisingly, with larger number of paths, the runtime increases. However, this is not a significant concern, since we find optimal solutions at `selectNumber` as low as 5. These numbers are representative of all applications and topologies we have considered.




Figure 15: Runtime and objective ratio as function of paths; optimality is achieved in most cases with as few as 5 paths per class

Path selection strategy: Fig. 16 shows the difference in the SOL-computed objective function based on the different path selection strategies as well as the original solution (where obtainable). This figure indicates that choosing an appropriate path selection strategy can provide substantial benefits for large topologies. In our experience, most problems lend themselves to a fairly obvious path selection strategy: those with need for load balancing should use `random` and those that are latency-sensitive should use `shortest`. If in doubt, however, both strategies can be attempted.

9. RELATED WORK

We already discussed the optimization applications that motivated SOL. Here we focus on other related work.

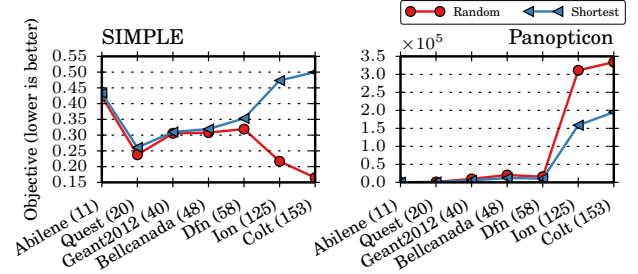


Figure 16: SOL objective values based on the chosen path selection strategy

Higher-layer abstractions for SDN: This work includes new programming languages (e.g., [32, 11]), testing and verification tools (e.g., [23]), semantics for network updates (e.g., [33]), compilers for rule generation (e.g., [22]), abstractions for handling control conflicts (e.g., [4]), and APIs for users to express requirements (e.g., [9]). These are orthogonal and do not address optimization in SDN applications, which is the focus of SOL.

Languages for optimization: There are several modeling frameworks such as AMPL [12], PyOpt [29], and PuLP [27] for expressing optimization tasks. However, these do not specifically simplify network optimization. SOL is a domain-specific library that operates at a higher level of semantics than these “wrappers”. SOL offers a path-based abstraction for writing network optimizations, exploits this structure to solve these optimizations quickly, and generates network device configurations that implement its solutions.

Network resource management: Merlin is a language for network resource management [36]. In terms of the specific applications that it can support, Merlin is restricted to using path predicates expressed as regular expressions. Our experiments suggest that SOL is three orders of magnitude faster than Merlin using the same underlying solvers. That said, Merlin’s “language-based” approach provides other capabilities (e.g., verified delegation) that SOL does not (try to) offer. Other works focus on traffic-steering optimization (e.g., [30, 6]). SOL offers a unifying abstraction that covers many network management applications.

10. CONCLUSION

Optimization is a core ingredient in designing many new SDN applications. Despite its broad utility, few efforts have attempted to make optimization more accessible to potential SDN application developers and network administrators. Our vision is a general and efficient high-level framework for expressing and solving complex network optimization tasks. We showed that SOL achieves both generality and efficiency via a path-centric optimization abstraction. This abstraction provides the generality to capture diverse applications, enables efficient solutions via simple path selection algorithms, and also simplifies SDN rule generation. We showed that

SOL can concisely capture optimization applications with diverse goals (traffic engineering, offloading, topology modification, service chaining, etc.), that SOL yields optimal or near-optimal solutions, and that SOL does so with far better online performance than handcrafted optimization formulations. We thus believe that SOL can significantly lower the barrier of entry and dramatically reduce development effort for novel SDN network optimization applications.

11. REFERENCES

- [1] R. K. Ahuja, T. L. Magnanti, and J. B. Orlin. *Network flows: theory, algorithms, and applications*. Prentice hall, 1993.
- [2] M. Al-Fares, A. Loukissas, and A. Vahdat. A scalable, commodity data center network architecture. In *ACM SIGCOMM Computer Communication Review*, volume 38, pages 63–74, 2008.
- [3] Anonymized. SOL manual, January 2015.
- [4] A. AuYoung, S. Banerjee, J. Lee, J. C. Mogul, J. Mudigonda, L. Popa, P. Sharma, and Y. Turner. Corybantic: Towards the modular composition of SDN control programs. In *ACM HotNets*, 2013.
- [5] A. Bremner-Barr, Y. Harchol, D. Hay, and Y. Koral. Deep packet inspection as a service. In *ACM CoNEXT*, pages 271–282, 2014.
- [6] Z. Cao, M. Kodialam, and T. Lakshman. Traffic steering in software defined networks: planning and online routing. In *ACM SIGCOMM Workshop on Distributed Cloud Computing*, pages 65–70, 2014.
- [7] M. Charikar, Y. Naamad, J. Rexford, and K. Zou. Multi-Commodity Flow with In-Network Processing. Manuscript, www.cs.princeton.edu/~jrex/papers/mopt14.pdf.
- [8] E. Danna, S. Mandal, and A. Singh. A practical algorithm for balancing the max-min fairness and throughput objectives in traffic engineering. In *IEEE Conference on Computer Communications*, pages 846–854, 2012.
- [9] A. Ferguson, A. Guha, C. Liang, R. Fonseca, and S. Krishnamurthi. Participatory networking: An API for application control of SDNs. In *ACM SIGCOMM*, August 2013.
- [10] B. Fortz and M. Thorup. Internet traffic engineering by optimizing OSPF weights. In *IEEE Conference on Computer Communications*, volume 2, 2000.
- [11] N. Foster, R. Harrison, M. J. Freedman, C. Monsanto, J. Rexford, A. Story, and D. Walker. Frenetic: A network programming language. In *ACM SIGPLAN Notices*, volume 46, pages 279–291, 2011.
- [12] R. Fourer, D. M. Gay, and B. W. Kernighan. *AMPL: A mathematical programming language*. AT&T Bell Laboratories Murray Hill, 1987.
- [13] A. Gember, A. Krishnamurthy, S. S. John, R. Grandl, X. Gao, A. Anand, T. Benson, V. Sekar, and A. Akella. Stratos: A network-aware orchestration layer for virtual middleboxes in clouds. *arXiv preprint arXiv:1305.0209*, 2013.
- [14] A. Gember-Jacobson, R. Viswanathan, C. Prakash, R. Grandl, J. Khalid, S. Das, and A. Akella. OpenNF: Enabling innovation in network function control. In *ACM SIGCOMM*, 2014.
- [15] G. Gibb, H. Zeng, and N. McKeown. Outsourcing network functionality. In *ACM SIGCOMM Workshop on Hot Topics in Software Defined Networking*, 2012.
- [16] B. Heller, S. Seetharaman, P. Mahadevan, Y. Yakoumis, P. Sharma, S. Banerjee, and N. McKeown. ElasticTree: Saving energy in data center networks. In *USENIX Symposium on Networked Systems Design and Implementation*, pages 19–21, 2010.
- [17] V. Heorhiadi, S. K. Fayaz, M. K. Reiter, and V. Sekar. SNIPS: A software-defined approach for scaling intrusion prevention systems via offloading. In *10th International Conference on Information Systems Security*, Dec. 2014.
- [18] V. Heorhiadi, M. K. Reiter, and V. Sekar. New opportunities for load balancing in network-wide intrusion detection systems. In *ACM CoNEXT*, 2012.
- [19] C.-Y. Hong, S. Kandula, R. Mahajan, M. Zhang, V. Gill, M. Nanduri, and R. Wattenhofer. Achieving high utilization with software-driven WAN. In *ACM SIGCOMM*, pages 15–26, 2013.
- [20] S. Jain, A. Kumar, S. Mandal, J. Ong, L. Poutievski, A. Singh, S. Venkata, J. Wanderer, J. Zhou, M. Zhu, et al. B4: Experience with a globally-deployed software defined WAN. In *ACM SIGCOMM*, pages 3–14, 2013.
- [21] X. Jin, L. E. Li, L. Vanbever, and J. Rexford. Softcell: Scalable and flexible cellular core network architecture. In *ACM CoNEXT*, 2013.
- [22] N. Kang, Z. Liu, J. Rexford, and D. Walker. Optimizing the one big switch abstraction in software-defined networks. In *ACM CoNEXT*, pages 13–24, 2013.
- [23] P. Kazemian, G. Varghese, and N. McKeown. Header space analysis: Static checking for networks. 2012.
- [24] S. Knight, H. Nguyen, N. Falkner, R. Bowden, and M. Roughan. The internet topology zoo. *IEEE Journal on Selected Areas in Communications*, 29(9):1765–1775, October 2011.
- [25] D. Levin, M. Canini, S. Schmid, F. Schaffert, A. Feldmann, et al. Panopticon: Reaping the benefits of incremental sdn deployment in enterprise networks. In *USENIX Annual Technical Conference*, 2014.
- [26] X. Liu, S. Mohanraj, M. Pioro, and D. Medhi. Multipath routing from a traffic engineering perspective: How beneficial is it? In *IEEE International Conference on Network Protocols*, pages 143–154, 2014.
- [27] S. Mitchell, M. O’Sullivan, and I. Dunning. Pulp: a

- linear programming toolkit for python, 2011.
- [28] Network functions virtualisation – introductory white paper. http://portal.etsi.org/NFV/NFV_White_Paper.pdf.
 - [29] R. E. Perez, P. W. Jansen, and J. R. R. A. Martins. pyOpt: A Python-based object-oriented framework for nonlinear constrained optimization. *Structures and Multidisciplinary Optimization*, 45(1):101–118, 2012.
 - [30] Z. A. Qazi, C.-C. Tu, L. Chiang, R. Miao, V. Sekar, and M. Yu. SIMPLE-fying middlebox policy enforcement using SDN. In *ACM SIGCOMM*, 2013.
 - [31] S. Raza, G. Huang, C.-N. Chuah, S. Seetharaman, and J. P. Singh. Measurouting: a framework for routing assisted traffic monitoring. *ACM/IEEE Transactions on Networking*, 20(1):45–56, 2012.
 - [32] J. Reich, C. Monsanto, N. Foster, J. Rexford, and D. Walker. Modular SDN programming with pyretic. *login: Magazine*, 38(5):128–134, 2013.
 - [33] M. Reitblatt, N. Foster, J. Rexford, C. Schlesinger, and D. Walker. Abstractions for network update. In *ACM SIGCOMM*, 2012.
 - [34] M. Roughan. Simplifying the synthesis of internet traffic matrices. *ACM SIGCOMM Computer Communication Review*, 35, 2005.
 - [35] J. Sherry, S. Hasan, C. Scott, A. Krishnamurthy, S. Ratnasamy, and V. Sekar. Making middleboxes someone else’s problem: Network processing as a cloud service. In *ACM SIGCOMM*, 2012.
 - [36] R. Soule, S. Basu, P. J. Marandi, F. Pedone, R. Kleinberg, E. G. Sirer, and N. Foster. Merlin: A language for provisioning network resources. In *ACM CoNEXT*, 2014.
 - [37] N. Vasić, P. Bhurat, D. Novaković, M. Canini, S. Shekhar, and D. Kostić. Identifying and using energy-critical paths. In *ACM CoNEXT*, 2011.
 - [38] R. Wang, D. Butnariu, and J. Rexford. Openflow-based server load balancing gone wild. In *Proc. Hot-ICE*, 2011.